COUNTING THE SPANNING TREES OF THE 3-CUBE USING EDGE SLIDES

CHRISTOPHER TUFFLEY

ABSTRACT. We give a direct combinatorial proof of the known fact that the 3-cube has 384 spanning trees, using an "edge slide" operation on spanning trees. This gives an answer in the case $n = 3$ to a question implicitly raised by Stanley. Our argument gives a bijective proof of the $n = 3$ case of a weighted count of the spanning trees of the n-cube due to Martin and Reiner, and we discuss the possibilities and difficulties of extending our approach to $n \geq 4$.

1. INTRODUCTION

The n-cube is the graph Q_n whose vertices are the subsets of the set $[n] = \{1, 2, \ldots, n\}$, with an edge between S and R if they differ by the addition or removal of precisely one element. The 3-cube is then the familiar graph shown in Figure 1, whose edges and vertices form the edges and vertices of an ordinary cube or die.

The number of spanning trees of the n-cube is known via Kirchoff's Matrix-Tree Theorem (see for example Stanley [5]) to be

$$|\mathrm{Tree}(Q_n)| = 2^{2^n-n-1} \prod_{k=1}^{n} k^{\binom{n}{k}} = \prod_{\substack{S\subseteq[n]\\|S|\geq 2}} 2|S|. \tag{1.1}$$

However, according to Stanley [5, p. 62] a direct combinatorial proof of this formula is not known. This situation is in stark contrast with that of a second infinite family of graphs, the complete graphs: the spanning trees of K_n may be counted not only via the Matrix-Tree Theorem, but also bijectively, using the Prüfer Code [4]. Indeed, there are several known proofs that K_n has n^{n-2} spanning trees — see for example Moon [3].

The purpose of this note is to give a direct combinatorial proof that the 3-cube has $2^4 \cdot 2^3 \cdot 3 = 384$ spanning trees, and thereby answer in the case $n = 3$ the question implicitly raised by Stanley's comment. We will do this by defining and using "edge slide" moves on trees to break the set of trees into families that are readily counted. Edge slide moves may also be defined for spanning trees of higher dimensional cubes, and we will briefly discuss the possibilities and difficulties of extending our approach to $n \geq 4$.

Our approach is motivated by the following refinement of (1.1), due to Martin and Reiner [2, Thm. 3]. Again using the Matrix-Tree Theorem, they prove a weighted count of the spanning trees of Q_n in terms of variables $q_1, \ldots, q_n$ and $x_1, \ldots, x_n$, namely

$$\sum_{T\in\mathrm{Tree}(Q_n)} q^{\mathrm{dir}(T)} x^{\mathrm{dd}(T)} = q_1 \cdots q_n \prod_{\substack{S\subseteq[n]\\|S|\geq 2}} \sum_{i\in S} q_i(x_i^{-1} + x_i). \tag{1.2}$$

Each term on the right-hand side corresponds to a tree, and is obtained by choosing, for each subset S of $[n]$ of size at least 2, an element i of S and a sign $\mu \in \{\pm 1\}$. We call such a series

2010 *Mathematics Subject Classification.* 05C30 (05C05).

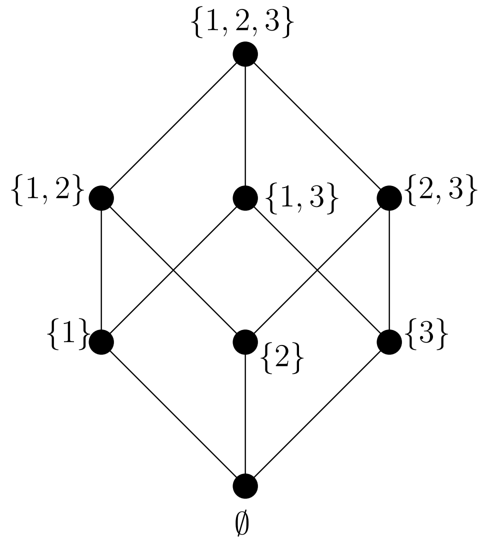

FIGURE 1. The three-cube.

of choices a *signed section* of $\mathcal{P}^n_{\geq 2} = \{S \in \mathcal{P}([n]) : |S| \geq 2\}$. Our combinatorial proof of (1.1) for $n = 3$ may be used to construct a weight preserving bijection between the spanning trees of Q_3 and the signed sections of $\mathcal{P}^3_{\geq 2}$, giving a combinatorial proof of the $n = 3$ case of (1.2).

2. Martin and Reiner's weighted count

Before proceeding we describe the weights used in Martin and Reiner's formula (1.2). The weight of a tree consists of two factors, the *direction monomial* $q^{\mathrm{dir}}(T)$ and the *decoupled degree monomial* $x^{\mathrm{dd}(T)}$, defined below in Section 2.1. In Section 2.2 we prove an alternate formulation of the decoupled degree monomial, which we will use in what follows.

2.1. **The direction and decoupled degree monomials.** Recall that we regard Q_n as the graph with vertex set the power set of $[n]$, with an edge between subsets S and R if they differ by the addition or deletion of a single element. If S and R differ by the addition or deletion of $i \in [n]$ we say that the edge $e = (S, R)$ is *in the direction* i, and write $\mathrm{dir}(e) = i$. With this notation, the direction monomial of a spanning tree T of Q_n is

$$q^{\mathrm{dir}(T)} = \prod_{e \in E(T)} q_{\mathrm{dir}(e)}.$$

To define the decoupled degree monomial, given $S \in [n]$ let $x_S = \prod_{i \in S} x_i$. Then

$$\begin{aligned} x^{\mathrm{dd}(T)} &= \prod_{S \subseteq [n]} \left(\frac{x_S}{x_{[n] \setminus S}} \right)^{\frac{1}{2} \deg_T(S)} \\ &= \prod_{(S,R) \in E(T)} \frac{x_S x_R}{x_{[n]}}. \end{aligned} \tag{2.1}$$

2.2. **An alternate formulation of the decoupled degree monomial.** Note that if the edge (S, R) of T is in the direction i then

$$\frac{x_S x_R}{x_{[n]}} = x_1^{\epsilon_1} \cdots x_{i-1}^{\epsilon_{i-1}} x_{i+1}^{\epsilon_{i+1}} \cdots x_n^{\epsilon_n}, \tag{2.2}$$

where for $j \neq i$,

$$\epsilon_j = \begin{cases} +1 & j \in S, \\ -1 & j \notin S. \end{cases} \tag{2.3}$$

Thus each edge of T in direction i contributes a factor of x_j or x_j^{-1} to $x^{\mathrm{dd}(T)}$ for each j *not* equal to i. The goal of this section is to show that the decoupled degree monomial can be re-expressed as a product over the edges e of T of $x_{\mathrm{dir}(e)}^{\pm 1}$, where the signs are determined by canonically orienting the edges of T.

The determine the signs orient all edges of Q_n "upwards", i.e. in the direction of increasing cardinality. Root each spanning tree T at the empty set, and orient each edge of T towards the root. Each edge e of T now has two orientations, one from the cube and one from the tree, and we let $\mu(e) = +1$ if the two orientations agree, and $\mu(e) = -1$ if the orientations disagree. So if $\mu(e)$ is $+1$ then to get to the root e must be crossed in the upwards direction relative to the cube, and if $\mu(e)$ is -1 then e must be crossed in the downwards direction to get to the root. With these signs we have

Lemma 2.1.

$$x^{\mathrm{dd}(T)} = x_1 x_2 \cdots x_n \prod_{e \in E(T)} x_{\mathrm{dir}(e)}^{\mu(e)}. \tag{2.4}$$

Proof. Delete all edges of T in the direction i. If there are k of them, this divides T into $k+1$ connected components, and $1 \leq j \leq k$ of them will be "upstairs" (vertices containing i), and the remaining $k+1-j$ will be downstairs. The part of T upstairs has 2^{n-1} vertices and Euler characteristic j (since its zeroth homology has rank j, and all higher homology groups are zero), so there are $2^{n-1} - j$ edges of T upstairs, and $2^{n-1} + j - k - 1$ downstairs. By (2.2) and (2.3), the degree of x_i in $x^{\mathrm{dd}(T)}$ is the number of edges in T upstairs minus the number of edges downstairs, or $k+1-2j$. Now each connected component of T upstairs must be adjacent to a unique downward edge in direction i, so j of the edges in direction i point down and the remaining $k-j$ point up. Thus the exponent of x_i in (2.4) is also $(k-j) - j + 1 = k+1-2j$, so (2.1) and (2.4) agree. □

Note that the factors $x_1, x_2, \ldots, x_n$ in front of the product in (2.4) are necessarily canceled by factors inside the product, since each spanning tree must have at least one downward edge in each direction. Thus if there are k_i edges in direction i, then the degree of x_i in $x^{\mathrm{dd}(T)}$ has the opposite parity to k_i and lies between $1 - k_i$ and $k_i - 1$.

3. Edge slides for the three-cube

3.1. **Definition and existence.** For each $i \in \{1, 2, 3\}$ let F_i^+ and F_i^- be the "upper" and "lower" faces of Q_3 with respect to direction i, in other words, the subgraphs induced by the vertices that respectively do and do not contain i. There is an obvious automorphism of Q_3 induced by a reflection that exchanges F_i^+ and F_i^-, and we will denote this automorphism by σ_i. In terms of the symmetric difference $\ominus$ this map is given by

$$\sigma_i(S) = S \ominus \{i\}.$$

Let T be a spanning tree of Q_3, and let e be an edge of T in a direction $j \neq i$ such that T does not also contain $\sigma_i(e)$. We will say that e is *i-slidable* or *slidable in direction i* if deleting e from T and replacing it with $\sigma_i(e)$ yields a second spanning tree T'. We may think of this operation as "sliding" e across a face of the cube to get a second spanning tree, as shown in

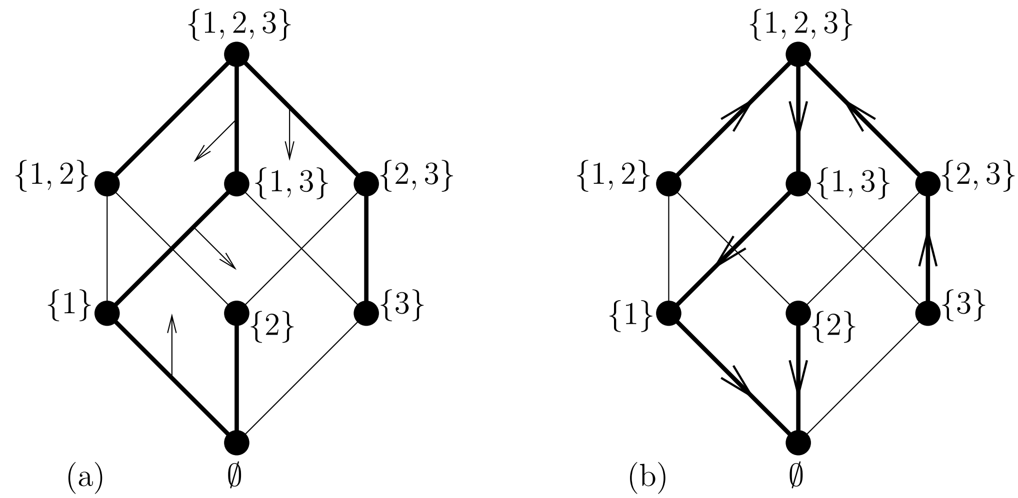

FIGURE 2. (a) The slidable edges in a certain spanning tree T of the three-cube. The edge $(\emptyset, \{1\})$ may be slid up in direction 2; the edge $(\{1\}, \{1,3\})$ may be slid down in direction 1; the edge $(\{1,3\}, \{1,2,3\})$ may be slid down in direction 3; and the edge $(\{2,3\}, \{1,2,3\})$ may be slid down in direction 2. (b) The oriented edges of T (see Section 3.2). The edges $(\{3\}, \{2,3\})$, $(\{2,3\}, \{1,2,3\})$ and $(\{1,2\}, \{1,2,3\})$ are upward edges and the rest are downward.

Figure 2. We will say that an edge slide is "upward" or "downward" according to whether it moves an edge from F_i^- to F_i^+, or from F_i^+ to F_i^-.

Adding any edge of Q_3 to T necessarily creates a cycle, so if $e \in T$ is i-slidable the cycle created by adding $\sigma_i(e)$ to T must be broken by deleting e. This cycle must therefore contain both e and $\sigma_i(e)$, and hence at least two distinct edges in direction i. It follows that any i-slidable edge must lie on the path between two edges in direction i. The following lemma shows that a minimal path joining two edges of T in direction i contains exactly one i-slidable edge, and has as a consequence the fact that a spanning tree of Q_3 has exactly four possible edge slides.

Lemma 3.1. *Let T be a spanning tree of Q_3, and let e_1, e_2 be edges of T in direction i. If the path P from e_1 to e_2 in T does not meet any other edge of T in direction i, then P contains a unique i-slidable edge e. Moreover, if T' is the result of sliding e in direction i, the slid edge $\sigma_i(e)$ is the unique i-slidable edge on the path from e_1 to e_2 in T'.*

Remark 3.2. In Section 5 we will see that we get existence but not uniqueness for spanning trees of Q_n, $n \geq 4$.

Proof. The lemma is proved by breaking it into cases according to the length of P. The path P lies in a face of Q_3, which is a square, and so has length at most three; we will treat only the case where P has length exactly three, as the remaining cases are similar but easier.

When P has length three it may be drawn as in Figure 3(a), in which the solid edges belong to T and P is the path (v_0, v_1, v_2, v_3). Consider the vertices $\sigma_i(v_1)$ and $\sigma_i(v_2)$. They must belong to T; but, since neither edge $(v_1, \sigma_i(v_1))$ nor $(v_2, \sigma_i(v_2))$ does, two of the three edges $(\sigma_i(v_0), \sigma_i(v_1))$, $(\sigma_i(v_1), \sigma_i(v_2))$ and $(\sigma_i(v_2), \sigma_i(v_3))$ must instead. It follows that T must be given by one of the three graphs shown in Figures 3(b)–(d). In each case, if $f = (\sigma_i(v_i), \sigma_i(v_{i+1}))$ is

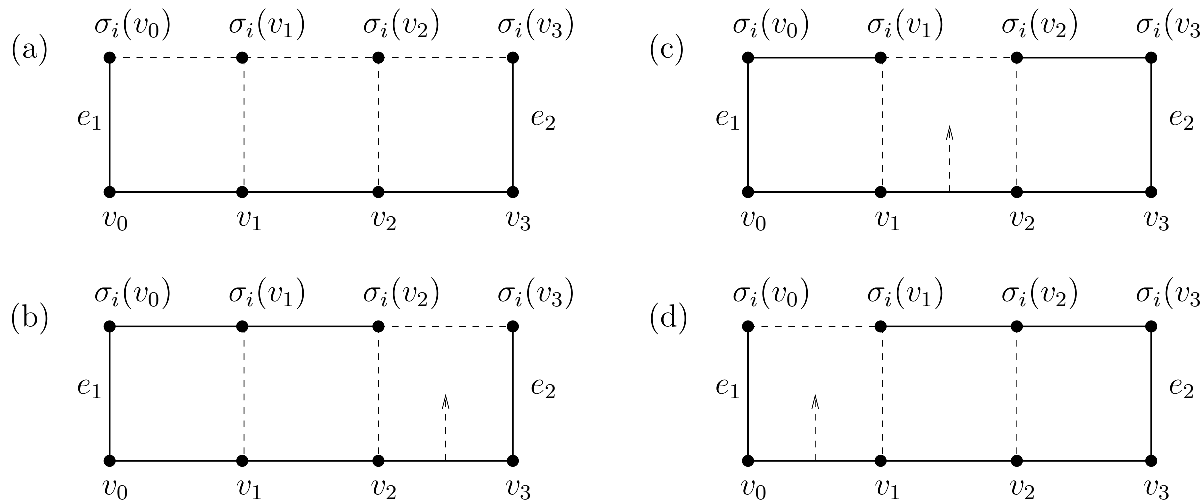

FIGURE 3. The case where P has length 3 in the proof of Lemma 3.1.

the dashed edge that does not belong to T, then $e = \sigma_i(f)$ is the unique i-slidable edge on P, and $f = \sigma_i(e)$ is the unique i-slidable edge on the path from e_1 to e_2 in T'. □

Lemma 3.3. *Let T be a spanning tree of Q_3, and suppose that T has u_i upward and d_i downward edges in direction i, for a total of $u_i + d_i = k_i$ edges in direction i. Then T has precisely $k_i - 1$ edges that may be slid in direction i, and of these u_i may be slid downwards, and the remaining $d_i - 1$ may be slid upwards.*

Proof. By Lemma 3.1 the i-slidable edges of T must totally disconnect the i-edges of T; since there are k_i edges in direction i at least $k_i - 1$ edges are required to totally disconnect them. So there are at least $k_i - 1$ i-slidable edges.

To bound the number of upward and downward i-slides from below we delete all edges of T in direction i, and consider the upper and lower faces separately. Deleting all the i-edges of T divides T into $k_i + 1$ connected components, and as seen in the proof of Lemma 2.1, a total of d_i of these components lie in F_i^+, with the remaining $u_i + 1$ in F_i^-. Now delete the i-slidable edges of T. Since this totally disconnects the i-edges of T it must further divide the d_i components upstairs into at least k_i components, requiring at least $k_i - d_i = u_i$ slidable edges upstairs; and similarly it must further divide the $u_i + 1$ components downstairs into at least k_i components, requiring at least $k_i - (u_i + 1) = d_i - 1$ slidable edges downstairs.

We now use the uniqueness clause of Lemma 3.1 to show that there can be no more than $k_i - 1$ edges that may be slid in direction i. Let $E = \{e_1, e_2, \ldots, e_{k-1}\}$ be a set of i-slidable edges that totally disconnect the i-edges of T, and let e be any i-slidable edge of T. Then e must lie on a path P in T from one i-edge of T to another, and we may choose P so that it does not meet any other i-edge of T. Then e is the unique i-slidable edge on P; but on the other hand, P must also cross an edge in E, since these totally disconnect the i-edges of T. It follows that $e = e_j$ for some j, so T has exactly $k_i - 1$ edges that may be slid in direction i. □

Corollary 3.4. *A spanning tree T of Q_3 has precisely four possible edge slides.*

Proof. This is immediate from Lemma 3.3 and the fact that T has seven edges, with at least one edge in each of the three directions. □

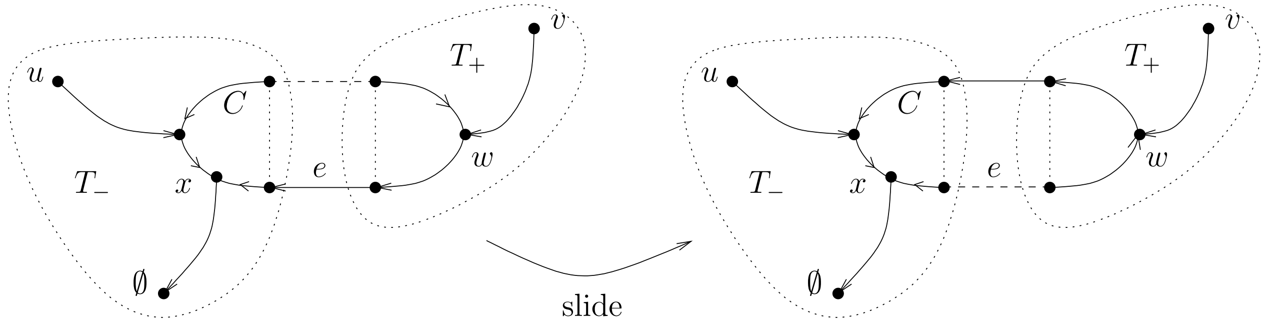

FIGURE 4. The effect of an edge-slide on orientations. Only those edges belonging to $C \cap T_+$ are reversed.

3.2. **Effect on orientations.** Clearly, the effect of an edge slide in direction i on the decoupled degree monomial is to multiply it by $x_i^{\pm 2}$: by x_i^2 if the edge is slid up, and x_i^{-2} if the edge is slid down. Thus, the number of upward edges in direction i must change by ± 1, and there can be no net change in the number of upward edges in other directions. The following lemma may be used to show that in the case of the 3-cube this occurs through a reversal of orientation of exactly one edge in direction i.

Lemma 3.5. *Given an i-slidable edge e of a spanning tree T, let C be the cycle created by adding $\sigma_i(e)$ to T; and let T_- be the component of $T - e$ that contains the root, and T_+ the component that does not. Then sliding e in direction i reverses the orientation of precisely those edges which belong to both C and T_+. The orientation of $\sigma_i(e)$ co-incides with that of e.*

Proof. Let T' be the tree resulting from the edge slide, and refer to Figure 4. If u is a vertex belonging to T_- then the paths from u to the root are the same in T and T', so edges in T_- have identical orientations in T and T'. If v is a vertex lying in T_+ then the path from v to the root may be expressed in the form PQR, where

- P is the path from v to the closest vertex w lying on C;
- Q is the path in C from w to x that crosses e, where x is the closest point on C to the root;
- R is the path from x to the root.

Then the path from v to the root in T' has the the form $PQ'R$, where Q' is the path in C from w to x that crosses $\sigma_i(e)$. It follows that edges of T_+ that do not lie on C are unchanged in orientation, and by considering the cases where v is the vertex of e or $\sigma_i(e)$ in T_+ we see that edges in $C \cap T_+$ are reversed. □

In the case of the 3-cube, with notation as in Lemma 3.1 it is clear that C consists of e, e_1, e_2 and P together with $\sigma_i(P)$. Thus, sliding e in direction i reverses the orientation of exactly one edge in direction i, namely whichever of e_1 and e_2 lies in T_+. Since the orientation of e points from T_+ to T_-, the edge reversed by sliding e is whichever of e_1 and e_2 that e points *away* from.

3.3. **Independent slides.** We show that parallel edge-slides on a spanning tree of the 3-cube commute in the following sense.

Let T be a spanning tree, and let $S = \{e_1, \ldots, e_k\}$ be a set of i-slidable edges of T. Given a vector $\varepsilon = (\varepsilon_1, \ldots, \varepsilon_k) \in \{0,1\}^k$, let T_ε be the subgraph

$$T_\varepsilon = \big(T \setminus S\big) \cup \{\sigma_i^{\varepsilon_1}(e_1), \ldots, \sigma_i^{\varepsilon_k}(e_k)\}. \tag{3.1}$$

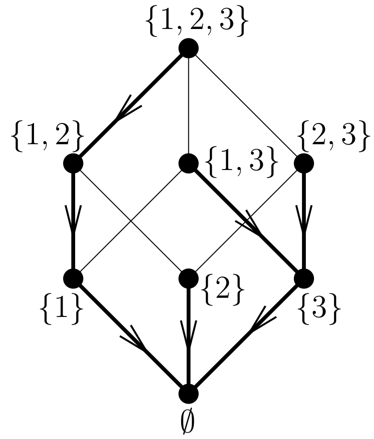

FIGURE 5. An upright spanning tree of Q_3. The associated section ϕ (see Section 4.3) is determined by $\phi(\{1,2,3\}) = 3$, $\phi(\{1,2\}) = 2$, $\phi(\{1,3\}) = 1$ and $\phi(\{2,3\}) = 2$.

Thus, T_ε is the subgraph of the cube obtained by choosing whether or not to slide each edge $e_j \in S$ according to the value of $\varepsilon_j \in \{0,1\}$. We will say S is an *independently slidable set* if T_ε is a spanning tree of the cube for all $\varepsilon \in \{0,1\}^k$.

In Lemma 3.6 we show that parallel edge slides in a spanning tree of the 3-cube are independent in the above sense, which will allow us to count the trees combinatorially.

Lemma 3.6. *The set $S_i(T)$ consisting of the i-slidable edges of a spanning tree T of the 3-cube is an independently slidable set.*

Proof. Referring again to Figure 3, we see that i-slidability of a given edge $e \in S_i(T)$ is a local property, depending only on the minimal cycle C obtained by adding $\sigma_i(e)$ to T. The edges of C are unaffected by i-slides of other edges in $S_i(T)$, and so e remains slidable regardless of how these other edges are slid. $\square$

4. Counting the spanning trees of the three-cube

4.1. **Strategy.** We will now use the results of the previous section to show that Q_3 has $2^4 \cdot 2^3 \cdot 3 = 384$ spanning trees. We will do this by constructing a projection from $\mathrm{Tree}(Q_3)$ onto a space of trees that are easily counted; showing that there are $2^3 \cdot 3$ trees in this family; and that each fibre of the projection has size 2^4.

For the target of the projection we define a spanning tree of Q_3 to be *upright* if all of its edges are oriented downwards. An example appears in Figure 5. We denote the set of upright trees of Q_3 by $\mathrm{UTree}(Q_3)$, and will count the trees by constructing a projection $\pi : \mathrm{Tree}(Q_3) \to \mathrm{UTree}(Q_3)$.

We will build the projection π up as a composition of retractions, and to this end we define a spanning tree to be *upright with respect to direction i* if all of its edges in direction i are oriented downwards. We write $\mathcal{D}_i(Q_3)$ for the set of trees that are upright with respect to direction i, and observe that

$$\mathrm{UTree}(Q_3) = \bigcap_{i=1}^{3} \mathcal{D}_i(Q_3).$$

We will proceed by constructing retractions $\pi_i : \mathrm{Tree}(Q_3) \to \mathcal{D}_i(Q_3)$, and then set $\pi = \pi_1 \circ \pi_2 \circ \pi_3$.

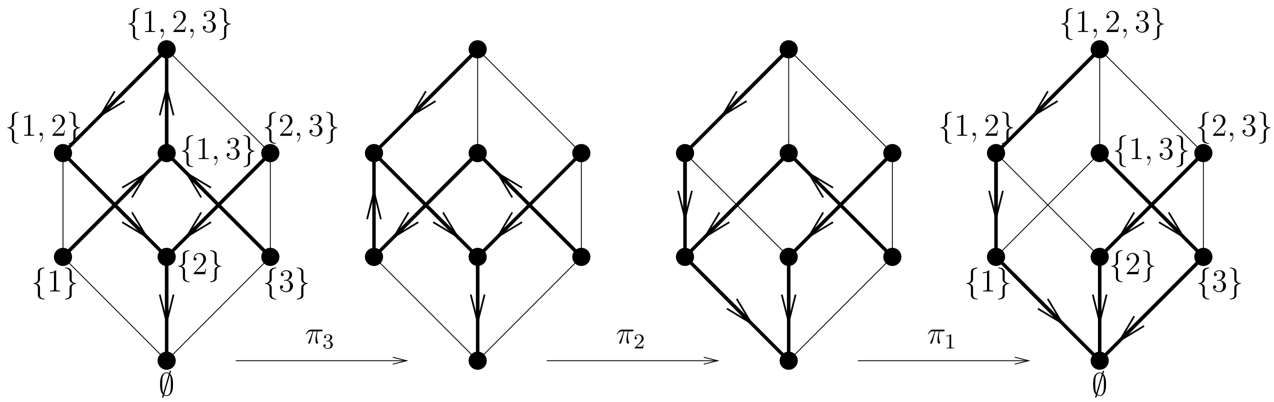

FIGURE 6. An example of the retraction $\pi = \pi_1 \circ \pi_2 \circ \pi_3$. By successively carrying out all possible downward edges slides in directions 3, 2, 1 in turn we arrive at the upright tree on the right.

4.2. **The retractions.** Given a spanning tree T of Q_3 set $S = S_i(T) = \{e_1, \ldots, e_{k_i-1}\}$ in equation (3.1), and let

$$\mathcal{E}_i(T) = \left\{T_\varepsilon | \varepsilon \in \{0,1\}^{|S_i(T)|}\right\}.$$

By Lemma 3.6 this set consists of 2^{k_i-1} spanning trees of Q_3, and we claim that these sets form a partition of $\mathrm{Tree}(Q_3)$. Indeed, suppose that $T' = T_\varepsilon \in \mathcal{E}_i(T)$. Then T_ε has the same number of i-edges as T, and hence the same number of i-slidable edges as T, by Lemma 3.3; thus we necessarily have $S_i(T_\varepsilon) = \{\sigma_i^{\varepsilon_1}(e_1), \ldots, \sigma_i^{\varepsilon_{k_i-1}}(e_{k_i-1})\}$, since these edges are i-slidable in T_ε. Then for $\delta \in \{0,1\}^{|S_i(T_\varepsilon)|}$ we have $(T_\varepsilon)_\delta = T_{\varepsilon+\delta}$, and so $\mathcal{E}_i(T_\varepsilon) = \mathcal{E}_i(T)$. The sets $\mathcal{E}_i(T)$ are therefore disjoint or equal, and since T necessarily belongs to $\mathcal{E}_i(T)$, they form a partition as claimed. We note in passing that these sets are the equivalence classes of the relation $\sim_i$ defined by $T_1 \sim_i T_2$ if and only if T_1 can be transformed into T_2 by a series of i-slides.

Now, for each spanning tree T of Q_3, carrying out all possible downward i-slides gives a unique choice of $T_{\tilde{\varepsilon}} \in \mathcal{E}_i(T)$ such that $\sigma_i^{\tilde{\varepsilon}_j}(e_j)$ lies in F_i^- for each $j = 1, \ldots, k_i - 1$. For such a tree only upward i-slides are possible, and so $T_{\tilde{\varepsilon}}$ can have only downward i-edges, by Lemma 3.3. Setting $\pi_i(T) = T_{\tilde{\varepsilon}}$ we therefore obtain a retraction $\pi_i : \mathrm{Tree}(Q_3) \to \mathcal{D}_i(Q_3)$. For each $T \in \mathcal{D}_i(Q_3)$ the fibre of this map is $\mathcal{E}_i(T)$, and so has cardinality 2^{k_i-1}, where k_i is the number of i-edges of T.

We now consider the composition $\pi = \pi_1 \circ \pi_2 \circ \pi_3$. For a spanning tree T of Q_3 the tree $\pi(T)$ is obtained by

(1) carrying out all possible downward edge slides in direction 3 in T to get $\pi_3(T)$; then
(2) carrying out all possible downward edge slides in direction 2 in $\pi_3(T)$, to get $\pi_2(\pi_3(T))$; then
(3) carrying out all possible downward edge slides in direction 1 in $\pi_2(\pi_3(T))$, to get $\pi_1(\pi_2(\pi_3(T))) = \pi(T)$.

An example appears in Figure 6. For this map we claim

Lemma 4.1. *The map $\pi = \pi_1 \circ \pi_2 \circ \pi_3$ is a retraction from* $\mathrm{Tree}(Q_3)$ *onto* $\mathrm{UTree}(Q_3)$*. For each tree $T \in \mathrm{UTree}(Q_3)$ the preimage of T contains exactly 2^4 trees.*

Proof. Recall from Section 3.2 that an i-slide has no net effect on the number of upward edges in directions other than i. This implies that

$$\pi_i(\mathcal{D}_j(Q_3)) \subseteq \mathcal{D}_j(Q_3) \quad \text{for all } i \text{ and } j, \tag{4.1}$$

and that

$$\pi_i^{-1}(\mathcal{D}_j(Q_3)) \subseteq \mathcal{D}_j(Q_3) \quad \text{for } j \neq i. \tag{4.2}$$

The relation

$$\pi(\mathrm{Tree}(Q_3)) = \pi_1 \circ \pi_2 \circ \pi_3(\mathrm{Tree}(Q_3)) \subseteq \mathcal{D}_1(Q_3) \cap \mathcal{D}_2(Q_3) \cap \mathcal{D}_3(Q_3) = \mathrm{UTree}(Q_3)$$

now follows immediately from from equation (4.1) and the fact that $\pi_i(\mathrm{Tree}(Q_3)) = \mathcal{D}_i(Q_3)$. Moreover if T is upright then $\pi_i(T) = T$ for all i, so $\pi(T) = T$ and π is a retraction.

Suppose now that T is an upright tree with k_i edges in direction i for $i = 1, 2, 3$. Then the preimage of T under π_1 is $\mathcal{E}_1(T)$, which consists of 2^{k_1-1} trees, each with the same number of edges in each direction as T, and each lying in $\mathcal{D}_2(Q_3) \cap \mathcal{D}_3(Q_3)$, by equation (4.2). The preimage of each tree $T_\varepsilon \in \mathcal{E}_1(T)$ under π_2 is then $\mathcal{E}_2(T_\varepsilon)$, which consists of 2^{k_2-1} spanning trees, each of which must belong to $\mathcal{D}_3(Q_3)$ and have k_3 edges in direction 3, and pulling each such tree $(T_\varepsilon)_\delta$ back under π_3 we get $\mathcal{E}_3((T_\varepsilon)_\delta)$, which consists of 2^{k_3-1} spanning trees. We therefore get a total of

$$2^{k_1-1} \times 2^{k_2-1} \times 2^{k_3-1} = 2^{k_1+k_2+k_3-3} = 2^4$$

trees in $\pi^{-1}(T)$. $\square$

4.3. **The number of upright trees.** Given an upright spanning tree T, the first edge on the path from a non-root vertex S to the root must be in a direction $\phi_T(S) = i$ belonging to S. To each upright tree we may therefore associate a function $\phi_T : \mathcal{P}([3]) \setminus \{\emptyset\} \to [3]$ such that $\phi_T(S) \in S$ for all S. We will say that such a function is a *section* of $\mathcal{P}^3_{\geq 1} = \mathcal{P}([3]) \setminus \{\emptyset\}$.

Conversely, given a section ϕ of $\mathcal{P}^3_{\geq 1}$, let $T(\phi)$ be the subgraph of Q_3 containing all eight vertices and the seven edges

$$\Big\{\{S, S \setminus \{\phi(S)\}\} : S \in \mathcal{P}^3_{\geq 1}\Big\}.$$

At each vertex S of Q_3 other than $\emptyset$ there is a unique edge in $T(\phi)$ connecting S to a vertex of cardinality $|S| - 1$, and following such edges one obtains a path in $T(\phi)$ from S to the root $\emptyset$. It follows that $T(\phi)$ is connected, and since it has seven edges and includes all eight vertices it must be a spanning tree. Moreover, these paths show that $T(\phi)$ is upright, and it's easily seen that it has associated section ϕ.

This all proves the following lemma, and then Corollary 4.3 follows immediately from Lemmas 4.1 and 4.2.

Lemma 4.2. *The upright spanning trees of Q_3 are in bijection with the sections of $\mathcal{P}^3_{\geq 1}$. Consequently there are $2^3 \times 3 = 24$ upright trees.*

Corollary 4.3. *The 3-cube has $2^4 \cdot 2^3 \cdot 3 = 384$ spanning trees.*

4.4. **A bijective count.** With some additional bookkeeping we may establish a bijection Φ between $\mathrm{Tree}(Q_3)$ and the set of *signed sections* of $\mathcal{P}^3_{\geq 2} = \{S \in \mathcal{P}([3]) : |S| \geq 2\}$. These are functions $\phi = \phi_d \times \phi_s : \mathcal{P}^3_{\geq 2} \to [3] \times \{\pm 1\}$ such that $\phi_d(S) \in S$ for all S. Given such a function ϕ we may define its weight to be

$$q^{\mathrm{dir}(\phi)} x^{\mathrm{sgn}(\phi)} = q_1 q_2 q_3 \prod_{S \in \mathcal{P}^3_{\geq 2}} q_{\phi_d(S)} x_{\phi_d(S)}^{\phi_s(S)},$$

and the bijection will be weight preserving in the sense that $q^{\mathrm{dir}(T)} x^{\mathrm{dd}(T)}$ will equal the weight of the associated section ϕ. This gives a bijective proof of the $n = 3$ case of the Martin-Reiner formula (1.2).

Given a spanning tree T of Q_3, the upright tree $\pi(T)$ has a canonical associated section $\phi_{\pi(T)}$ of $\mathcal{P}^3_{\geq 1}$. Restricting $\phi_{\pi(T)}$ to the sets of size two or more gives an (unsigned) section ϕ of $\mathcal{P}^3_{\geq 2}$, and this restriction completely determines $\phi_{\pi(T)}$ and hence $\pi(T)$. Moreover $q^{\mathrm{dir}(T)} = q^{\mathrm{dir}(\pi(T))} = q^{\mathrm{dir}(\phi)}$. We may therefore define $\Phi = \Phi_d \times \Phi_s$ so that

$$\Phi_d(T) = \phi_{\pi(T)}\Big|_{\mathcal{P}^3_{\geq 2}}.$$

It remains to define the signs, and we will do this by studying the way in which the edge slides taking T to $\pi(T)$ affect the orientations of the edges.

In Section 3.2 we saw that sliding an edge e of T in direction i reverses the orientation of exactly one edge in direction i. We claim that the reversed i-edge is the same for all trees $T_\varepsilon \in \mathcal{E}_i(T)$; more precisely, sliding whichever of $e, \sigma_i(e)$ belongs to T_ε reverses the orientation of the same i-edge f in all cases.

To see this, let C_e be the cycle formed by adding $\sigma_i(e)$ to T. Then as noted at the end of Section 3.2, C consists of two i-edges e_1 and e_2, and the path P joining them in T together with $\sigma_i(P)$, and the edge f reversed by sliding e is whichever of e_1 and e_2 that e points away from. Neither e nor $\sigma_i(e)$ lies on the corresponding cycle $C_{e'}$ for any other i-slidable edge $e' \in S_i(T)$, and therefore the orientation of e or $\sigma_i(e)$ cannot be affected by sliding e' or $\sigma_i(e')$, by Lemma 3.5. It follows that e and $\sigma_i(e)$ always point away from the same edge e_1 or e_2 regardless of how the other edges in $S_i(T)$ are slid, proving the claim.

To define $\Phi_s(T)$ at vertices where $\Phi_d(T)(S) = i$ we write $\pi = \pi_1 \circ \pi_2 \circ \pi_3$ in the form $\alpha \circ \pi_i \circ \beta$, where β is the composition of the retractions preceding π_i, and α the composition of those that follow. (Thus for example when $i = 3$, β is the identity and $\alpha = \pi_1 \circ \pi_2$.) The tree $\pi_i \circ \beta(T)$ is obtained from $\beta(T)$ by carrying out all possible downward i-slides, and we partition the i-edges of $\beta(T)$ into three sets P_i, N_i and Z_i:

- P_i consists of the i-edges of $\beta(T)$ that may be reversed by downward i-slides in $\beta(T)$ (and hence the i-edges that are reversed in obtaining $\pi_i \circ \beta(T)$ from $\beta(T)$);
- N_i consists of the i-edges of $\beta(T)$ that may be reversed by upward i-slides in $\beta(T)$; and
- Z_i consists of the unique i-edge of $\beta(T)$ that may not be reversed by an i-slide. This is the i-edge that belongs to the component containing the root when the i-slidable edges of $\beta(T)$ are deleted.

We note that this partition uniquely determines $\beta(T)$ from $\pi_i \circ \beta(T)$, by the discussion in the preceding paragraph.

We now keep track of this partition as $\pi_i \circ \beta(T)$ is transformed into $\pi(T)$ by α. During this transformation the i-edges may be slid in directions $j \neq i$, but we keep track of them through this movement to obtain a corresponding partition $\{P_i, N_i, Z_i\}$ of the i-edges of $\pi(T)$. There are now two possibilities:

(1) Z_i consists of the edge $\{\emptyset, \{i\}\}$. In this case we simply set $\Phi(T)(S) = (i, +1)$ if the first edge on the path in $\pi(T)$ to the root is an i-edge belonging to P_i, and $\Phi(T)(S) = (i, -1)$ if the first edge on the path in $\pi(T)$ to the root is an i-edge belonging to N_i.
(2) $Z_i = \{e\}$, for some i-edge $e \neq \{\emptyset, \{i\}\}$. In this case we modify the partition by swapping e and $\{\emptyset, \{i\}\}$, and then assign signs as in the previous case.

Example. We determine the signs associated with direction 3 for the tree appearing on Figure 6. In the leftmost tree the 3-slidable edges are $\{\{2\}, \{1, 2\}\}$, which may be slid up to reverse the orientation of the 3-edge $\{\{1, 2\}, \{1, 2, 3\}\}$, and $\{\{1, 3\}, \{1, 2, 3\}\}$, which may be slid down to reverse $\{\{1\}, \{1, 3\}\}$. So P_3 contains the edge $\{\{1\}, \{1, 3\}\}$ only, N_3 contains the edge

$\{\{1,2\},\{1,2,3\}\}$ only, and Z_3 contains the remaining 3-edge $\{\{2\},\{2,3\}\}$, which is not reversed by either 3-slide. Under the 2- and 1-slides the edge $\{\{1\},\{1,3\}\}$ belonging to P_3 is moved to $\{\emptyset,\{3\}\}$, putting us in case (2) above. The 3-edge belonging to Z_3 is still $\{\{2\},\{2,3\}\}$, so we assign the plus sign to the vertex $\{2,3\}$. The associated signed section in full is

$$\Phi(\{1,2\} = (2,+1), \qquad \Phi(\{1,3\}) = (1,+1)$$
$$\Phi(\{2,3\} = (3,+1), \qquad \Phi(\{1,2,3\}) = (3,-1).$$

It is clear by construction that the resulting map Φ is weight preserving, so it remains to check that T is in fact determined by the associated signed section $\Phi(T)$. The upright tree $\pi(T)$ may be recovered from $\Phi_d(T)$, and we partition the i-edges of $\pi(T)$ as $\{P_i, N_i, Z_i\}$ according to the signs. We may then carry out slides in directions 1, 2 and 3 in turn so as to reverse the orientations of the edges belonging to P_i for each i, keeping track of each partition as we do so. The only difficulty that can arise is if at the ith stage the i-edge that cannot be reversed belongs to P_i (this can only occur for $i \geq 2$). In that case we reverse the edge belonging to Z_i in its place, which has the effect of undoing the modification made to the partition in case (2) above.

5. Edge slides in higher dimensions

The definition of an edge slide in Section 3.1 was stated only for a spanning tree of the three-cube, but it applies just as well to a spanning tree of Q_n for any $n \geq 2$. In this section we explore the possibilities and difficulties of using edge slides to count the spanning trees of Q_n, for $n \geq 4$.

In the positive direction, we prove in Section 5.1 that a spanning tree with k_i edges in direction i always has at least $k_i - 1$ edges that may be slid in direction i, and that any spanning tree may be transformed into an upright tree by a sequence of downward edge slides only. Moreover, the argument of Section 4.3 may be used to show that the upright trees of Q_n are in bijection with the sections of $\mathcal{P}^n_{\geq 1}$, so that Q_n has a total of

$$\prod_{k=1}^{n} k^{\binom{n}{k}}$$

upright trees. Thus, we may hope to count the trees by constructing a retraction $\pi : \mathrm{Tree}(Q_n) \to \mathrm{UTree}(Q_n)$, analogous to that for $n = 3$ in Section 4.2.

However, our construction of the retraction $\pi : \mathrm{Tree}(Q_3) \to \mathrm{UTree}(Q_3)$ depended on the uniqueness clause of Lemma 3.1, and the independence of parallel edge slides of Lemma 3.6. In the negative direction we show by example in Section 5.2 that these results fail for $n \geq 4$. Thus, some additional ideas are required if such retractions are to be constructed.

A natural approach to further study of the edge slide operation is to look at the graph it generates, with the spanning trees of Q_n as the vertices. We conclude by defining this *edge slide graph* in Section 5.3, and briefly discussing the connected components of the edge slide graph of the three-cube.

5.1. **Existence.** As in the three-dimensional case, an i-slidable edge must lie on the path between two edges in direction i. We prove the following existence theorem, and deduce two corollaries.

Theorem 5.1. *Let T be a spanning tree of Q_n, and let e_1 and e_2 be edges of T in direction i. Then there is an i-slidable edge on the path from e_1 to e_2 in T.*

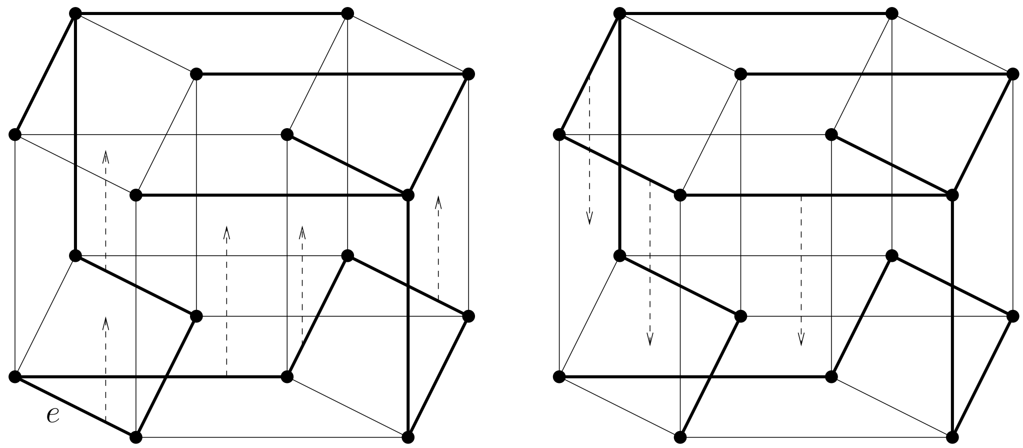

FIGURE 7. "Extra" edge slides when $n = 4$. The tree on the left has only two vertical edges, but there are five edges (indicated with dashed arrows) on the path joining them that may be slid vertically. The tree on the right shows the result of sliding the edge labelled e at left. This now has three vertically slidable edges on the path joining the two vertical edges.

Proof. Without loss of generality we may assume that the path P from e_1 to e_2 in P does not meet any other edge of T in direction i. Let $P = (v_0, v_1, \dots, v_m)$, where the vertices v_0 and v_m are incident with e_1 and e_2 respectively.

For each vertex v_j of P let $\phi(v_j)$ be the first vertex of P on the path from $\sigma_i(v_j)$ to v_j in T. Clearly, $\phi(v_0) = v_0$, and $\phi(v_m) = v_m$. Suppose that there is an edge $f = (v_\ell, v_{\ell+1})$ of P such that $\phi(v_\ell)$ and $\phi(v_{\ell+1})$ lie on *opposite* sides of f. Then adding $\sigma_i(f)$ to T creates a cycle that is broken by deleting f, so f is i-slidable.

If there is no such edge f then we may show by induction that $\phi(v_j) \in \{v_0, v_1, \dots, v_{j-1}\}$ for $1 \le j \le m$. But this contradicts the fact that $\phi(v_m) = v_m$, so there must be an i-slidable edge on P. □

Corollary 5.2. *Let T be a spanning tree of Q_n, and suppose that T has u_i upward and d_i downward edges in direction i, for a total of $u_i + d_i = k_i$ edges in direction i. Then T has at least $k_i - 1$ edges that may be slid in direction i, and of these at least u_i may be slid downwards, and at least $d_i - 1$ may be slid upwards.*

Proof. By Theorem 5.1 the i-slidable edges of T totally disconnect the i-edges of T. The result now follows by the argument of the first two paragraphs of the proof of Lemma 3.3. □

Corollary 5.3. *Given a spanning tree T of Q_n there is a sequence of downward edge slides that transforms T into an upright tree.*

Proof. If T has an upward edge in direction i then by Corollary 5.2 there is a downward edge slide in direction i. As discussed at the beginning of Section 3.2 the net effect of this slide is to decrease the number of upward i-edges by one, with no change to the number of upward edges in other directions. The total number of upward edges may therefore be decreased monotonically by a sequence of downward slides, and after $\sum u_i$ slides we will have an upright tree. □

5.2. **Counterexamples.** Our construction of the retraction $\pi : \mathrm{Tree}(Q_3) \to \mathrm{UTree}(Q_3)$ depended on the fact that a spanning tree of Q_3 with k_i edges in direction i has precisely $k_i - 1$

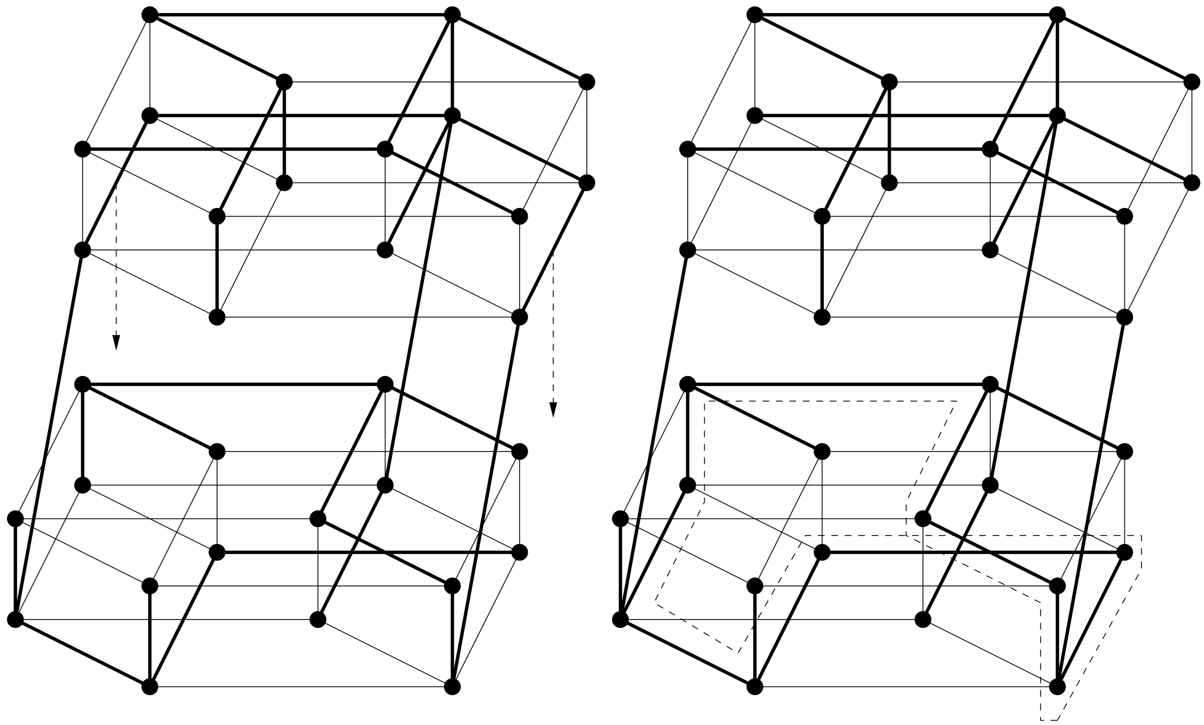

FIGURE 8. Parallel edge slides need not be independent, even when there are no more edge slides in that direction than expected. The figure on the left illustrates a spanning tree of the 5-cube, with some edges of the 5-cube omitted for clarity. The tree has three edges joining the upper and lower 4-cubes, and precisely two edges that may be slid from one 4-cube to the other. The figure on the right shows the result of sliding both. This contains a cycle, indicated by the dashed line which shadows it.

i-slidable edges, which may all be slid independently (in the sense that sliding any one of them has no effect on the slidability of the others). In Figures 7 and 8 we show by example that this fails for $n \geq 4$.

The tree on the left in Figure 7 has only two vertical edges, but five vertically slidable edges on the path joining them. When any one of these five edges is slid vertically the other four necessarily cease to be slidable, because the vertical edges are now disconnected downstairs. However, if the edge labelled e is slid two vertically slidable edges are created upstairs, as seen in the tree on the right.

Figure 8 shows that even when a spanning tree has precisely $k_i - 1$ i-slidable edges, the i-slides cannot necessarily be made independently. The tree on the left has three edges joining the upper and lower 4-cubes, and precisely two edges that may be slid from one 4-cube to the other. However, if both are slid the result is not a tree, as seen in the graph on the right.

To construct the retraction π it was necessary to choose an ordering of the directions 1, 2 and 3, but once this choice was made no further choices were necessary. The examples in this section show that for $n \geq 4$ it is also necessary to chose which edges to slide, and in what order to slide them. The difficulty in extending the ideas of this paper to higher dimensions lies in finding a way to make these choices systematically.

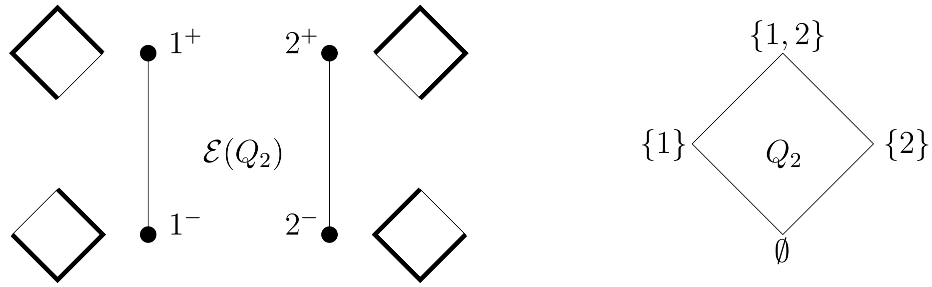

FIGURE 9. *(Left)* The edge slide graph of the 2-cube. There are four spanning trees, which are connected by edge slides in pairs, so the edge slide graph consists of two copies of K_2. The labels are a weight preserving bijection between the spanning trees and the signed sections of $\mathcal{P}^2_{\geq 2}$, when the vertices are labelled as shown on the right.

5.3. **The edge slide graph.** We define the *edge slide graph of* Q_n to be the graph $\mathcal{E}(Q_n)$ with vertices the spanning trees of Q_n, and an edge between trees T_1 and T_2 if they are related by a single edge slide. The graph $\mathcal{E}(Q_2)$ appears in Figure 9. The retraction $\pi : \mathrm{Tree}(Q_3) \to \mathrm{UTree}(Q_3)$ of Section 4.2 is based on a suitably chosen spanning forest of $\mathcal{E}(Q_3)$, so an understanding of $\mathcal{E}(Q_n)$ may help in constructing the necessary retractions. We conclude the paper by describing the connected components of this graph in the case $n = 3$.

Edge slides do not change the direction monomial $q^{\mathrm{dir}(T)} = q_1^{k_1} q_2^{k_2} q_3^{k_3}$, so if $q^{\mathrm{dir}(T_1)} \neq q^{\mathrm{dir}(T_2)}$ then T_1 and T_2 lie in different components. Thus, it suffices to understand the subgraphs $\mathcal{E}(k_1, k_2, k_3)$ consisting of the trees with direction monomial $q_1^{k_1} q_2^{k_2} q_3^{k_3}$. Moreover, any permutation of $\{1, 2, 3\}$ induces an automorphism of Q_3, and hence of $\mathcal{E}(Q_3)$, so we may consider the triple (k_1, k_2, k_3) up to permutation. We will refer to this triple as the *signature* of T, and up to permutation we find that there are three possible signatures, namely $(4, 2, 1)$, $(3, 3, 1)$ and $(3, 2, 2)$.

A spanning tree with $k_3 = 1$ consists of two spanning trees of Q_2, lying in F_3^- and F_3^+, joined by an edge in direction 3. No edge slide in direction 3 is possible, while the spanning trees of Q_2 in F_3^- and F_3^+ each have a single possible edge slide, and the edge joining them may be slid in either direction 1 or 2. It is easily seen that these four slides may be made independently, so each such tree belongs to a component of $\mathcal{E}(Q_3)$ isomorphic to Q_4. We get one such component for signature $(4, 2, 1)$, and two for $(3, 3, 1)$, characterised by the signature (k_1', k_2') of the spanning tree of F_3^-, which is invariant under edge slides in directions 1 and 2.

The subgraph $\mathcal{E}(3, 2, 2)$ is more interesting, as now edge slides in all three directions are possible. We find that there are four possible upright trees with signature $(3, 2, 2)$, occurring in two mirror image pairs, and that any two of these may be connected by a series of edge slides. Since any spanning tree may also be connected to an upright tree by a series of edge slides this implies that $\mathcal{E}(3, 2, 2)$ is connected. This gives us a connected component with 64 vertices, consisting of the 16×4 spanning trees associated with these four upright trees by π.

In total $\mathcal{E}(Q_3)$ has $6 \times 1 + 3 \times 2 = 12$ components isomorphic to Q_4, and three 64-vertex components isomorphic to $\mathcal{E}(3, 2, 2)$, accounting for all $24 \times 16 = 384$ spanning trees of Q_3. The structure of the 64-vertex component has been found by Lyndal Henden [1] as an undergraduate summer research project.

References

[1] L. Henden. The edge slide graph of the 3-cube. Submitted, 2011.

[2] J.L. Martin and V. Reiner. Factorizations of some weighted spanning tree enumerators. *J. Comb. Theory Ser. A*, 104(2):287–300, 2003.

[3] J.W. Moon. Various proofs of Cayley's formula for counting trees. In *A seminar on Graph Theory*, pages 70–78. Holt, Rinehart and Winston, 1967.

[4] H. Prüfer. Neuer Beweis eines Satzes über Permutationen. *Arch. Math. Phys.*, 27:742–744, 1918.

[5] R.P. Stanley. *Enumerative Combinatorics, Volume II.* Number 62 in Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1999.

Institute of Fundamental Sciences, Massey University, Private Bag 11 222, Palmerston North 4442, New Zealand

E-mail address: `c.tuffley@massey.ac.nz`